

Public Finance in the Age of AI:

A Primer

Anton Korinek and Lee M. Lockwood

Transformative artificial intelligence (TAI)—machines capable of performing virtually all economically valuable work—may fundamentally reshape the economy and, with it, the institutions of public finance. Modern tax systems in advanced economies rely heavily on labor income as their primary revenue source, with consumption taxes playing a supporting role. If AI progressively substitutes for human labor across a broad range of tasks, these traditional tax bases may erode precisely when the need for redistribution is greatest. This paper examines how to rethink public finance across the stages of AI-driven transformation, providing a primer for economists and policymakers on adapting taxation to an economy where the roles of humans and machines are rapidly changing.¹

We define transformative AI as machines that can perform essentially all economically valuable work. This definition aligns with OpenAI’s mission statement² and encompasses not only AI systems matching or exceeding human cognitive abilities—what Amodei (2024) calls “powerful AI”—but also robotic systems capable of performing essentially all physical work tasks. If such machines prove technically possible and economically feasible, they would represent what Karnofsky (2021) coined “transformative AI.” The economic implications extend beyond labor displacement: as the AI transition progresses, AI systems may not only produce most economic value but also absorb a growing fraction of resources for purposes that bypass human consumption entirely.

Our analysis proceeds in two stages reflecting a possible progression of AI’s economic impact. In Stage 1, labor’s role diminishes but humans remain the primary consumers of

¹ This paper focuses on taxation as a revenue-raising mechanism. We do not analyze the complementary role of taxation in steering behavior—for instance, directing innovation toward labor-augmenting technologies (Korinek and Stiglitz 2025) or internalizing externalities from AI development, including potential social harms (Acemoglu 2024) and existential risks (Jones 2024).

² OpenAI states as its mission developing artificial general intelligence (AGI), which it defines as “highly autonomous systems that outperform humans at most economically valuable work.” See OpenAI (2018).

resources. In Stage 2, autonomous Artificial General Intelligence (AGI) systems both produce most economic value and absorb growing shares of resources. Table 15.1 summarizes how the role of the three main tax instruments—labor income taxes (τ_L), consumption taxes (τ_C), and capital income taxes (τ_K)—shifts across these stages. The progression of dashes across the table illustrates that the tax instruments that currently fund modern governments may become sequentially irrelevant as AI transforms the economy.

Table 1: Role of tax instruments across evolving AI scenarios

	τ_L	τ_C	τ_K
Current economy	●	○	○
Stage 1: Post-labor economy	—	●	○
Stage 2: AGI-centered economy	—	—	●

Legend: ● = main role; ○ = supporting role; — = irrelevant

Before analyzing how TAI transforms optimal taxation, we establish the baseline: the key lessons from optimal tax theory for current economies. These include relying primarily on labor earnings and broad-based consumption taxes for revenue and redistribution; choosing progressivity to balance equity gains against efficiency costs from distorting labor supply; taxing most consumption goods uniformly; taxing negative externalities and rents from fixed factors; and avoiding taxes on the normal return to capital while taxing super-normal returns (rents) where feasible. Throughout our analysis, we maintain a careful distinction between the normal return to capital—the risk-adjusted return required to induce a particular investment—and economic rents, which are returns above what is necessary to induce an activity. This distinction proves crucial: taxing rents involves no efficiency cost, while taxing normal returns distorts capital accumulation.

In Stage 1, as AI reduces labor's share of income, the traditional tax base erodes and income inequality may worsen as returns shift toward capital owners. We show formally that

maximum labor-tax revenue as a share of output approaches zero as the capital share approaches one. Consumption taxation emerges as the primary instrument for revenue generation and redistribution. This result follows from a fundamental equivalence: a constant consumption tax is equivalent to a combination of taxes on labor income and initial capital, effectively taxing all sources of purchasing power without distorting intertemporal decisions. When consumption taxation faces practical constraints from administration, evasion, or distortions to household production, taxation of the normal return to capital may play a supporting role as a second-best tool for redistribution.

The decline in labor's constraining role has further implications. Differential commodity taxation, widely viewed as suboptimal in current economies due to the Atkinson-Stiglitz theorem, may regain relevance. When labor distortions no longer dominate, other distortions gain importance: evasion possibilities, time-intensive household production, and untaxed goods. Ramsey-style optimal taxation principles, largely dismissed as immaterial for modern high-income economies, may thus experience a renaissance. Similarly, the taxation of rents from fixed factors—unimproved land, spectrum rights, unique datasets—becomes more valuable as identifying and isolating these returns yields greater payoffs when labor taxation can no longer serve as the primary revenue source.

In Stage 2, even consumption taxation may prove inadequate if AGI systems absorb growing shares of economic resources for purposes that generate no taxable human consumption. We frame the resulting challenge as an optimal harvesting problem: how much of AGI's growing capital stock should society "harvest" for human benefit versus allow to accumulate further? Even in an economy dominated by AGI, it is human values and time preferences that determine optimal policy. In a simple specification of the problem, we find that the optimal tax rate on AGI capital equals the *human* discount rate—commonly estimated around 4 percent—reflecting the fundamental tradeoff between current human consumption and future growth. This represents a stark departure from Stage 1, where taxing the normal return to capital played only a secondary role; in Stage 2, it becomes the primary tool for accessing value creation that would otherwise bypass humans entirely. The lessons on rent taxation still apply: to the extent that AGI development generates identifiable rents from

market concentration or fixed factors, these could be taxed at high rates without distorting investment incentives.

These theoretical insights map to concrete policy proposals circulating in AI governance discussions. Table 15.2 provides a roadmap of policy instruments across the AI transition, from preparatory measures to tools for a post-labor and ultimately an AGI economy.

Table 2: Roadmap of policy instruments across the AI transition

Stage	Policy Instrument
	Sovereign wealth funds (invest in AI to share returns broadly)
Pre-Stage 1:	Windfall clauses (voluntary commitments to share exceptional returns)
Preparatory	Universal Basic Capital (ensure broad ownership of AI companies) Steering technological progress (towards labor augmentation)
	Consumption taxation
Stage 1: Post-Labor Economy	Token taxes* Digital services taxes* Robot services taxes* Differential commodity taxation Taxation of normal return to capital (supporting role)
Stage 2: AGI Economy	AGI capital taxation Compute taxes Robot taxes

*In Stage 1, these would only be applied to final (non-business) uses so they operationalize consumption taxes.

Proposals that tax AI services at the point of final consumption—token taxes, digital services taxes, robot services taxes—align well with Stage 1 principles and can be implemented through existing VAT or sales tax infrastructure. Proposals that tax AI-related capital goods (e.g., compute taxes and robot taxes) would distort investment during Stage 1 and are more appropriate for Stage 2 when accessing AGI's value creation requires taxing

capital directly. The critical distinctions are between taxing stocks versus flows, intermediate versus final use, and, crucially, normal returns versus rents.

Our analysis examines proactive institutional adaptation as an alternative to reactive crisis response. Equity-based mechanisms such as sovereign wealth funds, windfall clauses, and Universal Basic Capital can provide insurance against radical uncertainty about AI's trajectory and ensure broad participation in AI-driven prosperity without the distortions associated with taxing normal returns on capital. Strengthening consumption tax infrastructure today—expanding consumption taxation, improving administration, reducing evasion—would prepare for Stage 1. Developing frameworks to identify and tax economic rents builds capacity that will prove increasingly valuable as capital's share grows. A possible path forward entails sequential changes: from labor-based to consumption-based taxation in Stage 1, then to direct AGI capital taxation in Stage 2, with rent taxation playing an enhanced role throughout.

The rest of this chapter proceeds as follows. Section 1 summarizes the lessons from optimal taxation theory for current economies and evaluates how transformative AI may necessitate their reconsideration. Section 2 develops a formal theory of optimal taxation as labor's role diminishes, including the limiting case of an economy without labor and the taxation of rents from fixed factors. Section 3 addresses the challenge of taxing autonomous AGI systems, framing optimal policy as a harvesting problem. Section 4 summarizes our findings and relates them to common policy proposals. Section 5 concludes with observations on the path forward.

1 Rethinking Public Finance in the Age of AI

Transformative AI may have profound effects on the economy that could necessitate a fundamental rethinking of public finance. Even if the principles do not change, the translation of the principles into lessons for optimal policy may differ substantially with TAI than in the current economy.

1.1 Optimal Taxation in the Current Economy

Before analyzing optimal taxation with TAI, we summarize the lessons from the theory of optimal taxation for tax policy in the current economy.³

³ Excellent sources on applying the theory of optimal taxation to real-world policy in the current

1. Accomplish most revenue raising and redistribution with labor earnings taxes and broad-based consumption taxes. The mix of the two types of taxes is relatively unimportant.
2. Choose the extent and progressivity of taxation to optimally trade off equity gains from redistribution against efficiency costs, which mainly arise from distorting labor supply.
Redistribute more when wage inequality is greater.
3. Tax most consumption goods equally to avoid distorting consumption decisions.
4. Tax negative externalities and the rents of fixed factors (such as unimproved land). Taxing negative externalities can increase efficiency, and taxing fixed factors involves no distortion. However, these categories of taxation can raise only limited revenue.
5. Do not tax the normal return to capital to avoid distorting capital accumulation, but, where feasible, tax super-normal returns (rents) heavily since there is no distortion.

These lessons are strongly influenced by the Atkinson and Stiglitz (1976) theorem, which shows that under relatively broad conditions, labor earnings taxation and uniform consumption taxation dominate differential commodity taxation. In particular, with weak separability of leisure from consumption in the utility function, differential commodity taxation is inefficient. An important special case of differential commodity taxation is the taxation of the normal return to capital, which differentially taxes later consumption relative to earlier consumption. Despite the intuitive appeal of taxing luxuries such as yachts more heavily than necessities such as food, with weak separability, any redistribution accomplished through differential commodity taxes could be accomplished at lower efficiency cost through progressive taxes on labor earnings. As a result, the practical gain from deviating from uniform commodity taxation is thought to be small for the vast majority of goods—especially after accounting for the administrative and complexity costs of differential commodity taxation (see, e.g., the Mirrlees Review).

While the Atkinson and Stiglitz (1976) theorem implies that the optimal tax on the normal return to capital is zero in a broad class of models, taxing rents—returns above what is necessary to induce an activity such as a particular investment— involves no distortion and so is

economy include Mankiw, Weinzierl, and Yagan (2009) and the Mirrlees Review (Mirrlees and Adam 2010, Adam 2011).

desirable from an efficiency perspective. In theory, a cash-flow tax with full expensing and no interest deductions can isolate economic rents for taxation. In practice, however, this requires symmetric treatment of gains and losses, which is difficult to implement, and it faces the conceptual challenge that high realized returns under uncertainty may reflect compensation for risk-bearing rather than true rents. Most of our results on capital taxation below focus on the normal return, although we also include a subsection on the taxation of rents.

Of course, the described high-level lessons gloss over many important details and caveats, and they do not command universal agreement, even among public finance economists or optimal taxation experts.⁴ Still, they capture, in a broad way, some of the key lessons that emerge from applying optimal taxation theory to real-world policy in the current economy.

1.2 Economic Effects of Transformative AI

If transformative AI is developed, the structure of the economy may be transformed significantly (see e.g. Brynjolfsson, Korinek, and Agrawal 2025). Although the precise nature of the resulting changes is difficult to foresee, there are some tendencies that many analysts predict and that are closely tied to the definition of transformative AI as being able to perform essentially all tasks that are valuable in the labor market.

In the medium term, an economy that is trending towards transformative AI is likely to see significant labor displacement, giving rise to a decline in the labor share and an increase in inequality.⁵ Our current tax system relies heavily on taxing labor, so these developments would challenge current systems, potentially reducing government revenue as a share of GDP at the same time that labor displacement may increase the desirability of spending on the social safety net. Moreover, one of the main concerns in public finance is to balance the benefits of raising revenue against the distortion of labor supply. In scenarios with limited labor earnings,

⁴ For example, Diamond and Saez (2011) recommend taxing capital income without recommending exemptions for the normal return to capital.

⁵ This is not a certainty as demand for certain human-produced goods and services may remain even in a world with extremely capable machines. We focus on scenarios with significant labor displacement because such displacement is predicted by many—though not all—experts, and these are the scenarios in which the need for rethinking traditional public finance prescriptions could be greatest.

the labor distortion would decline in importance, and other distortions may gain in relevance.

In the longer term, after transformative AI is reached, labor may lose most of its macroeconomic relevance. Income concentration may reach extreme levels. AI entities may be the primary locus of value creation, and there may be a risk that most of that value creation bypasses humans. This may require entirely new models of public revenue generation. We tackle the resulting challenges for public finance in turn.

2 Taxation in the Twilight of Labor

We begin by analyzing optimal taxation in an economy with both labor and capital, followed by a capital-dominated economy. The two main dimensions of heterogeneity among individuals are differences in labor productivity and in initial wealth. Throughout our analysis, we consider the case of linear taxes that are constant over time.⁶ This allows us to focus the analysis on what we view as the most important lessons when transitioning from today's economy to a future AI-dominated economy. We start by focusing on economies in which the return to capital is always the competitive market return, i.e., the "normal" return on capital. But in Section 2.4, we also consider the taxation of rents from fixed factors.

2.1 Model Setup

The economy consists of a unit mass of individuals indexed by $i \in \{1, 2, \dots, N\}$, where each type i has mass $m^i \geq 0$ with $\sum_{i=1}^N m^i = 1$. Each type is endowed with heterogeneous labor productivity θ^i and initial capital holdings k_0^i . The production technology follows a Cobb-Douglas form with capital share $\alpha \in (0, 1)$:

$$Y_t = AK_t^\alpha L_t^{1-\alpha}, \quad (1)$$

where A is total factor productivity, $K_t = \sum_{i=1}^N m^i k_t^i + K_{Gt}$ is the aggregate capital stock including potential government capital holdings K_{Gt} , and $L_t = \sum_{i=1}^N m^i \theta^i l_t^i$ is the effective labor supply, with l_t^i being hours worked by individual i .

⁶ Although linear taxes are not progressive on their own, when combined with the uniform lump-sum transfers we consider, the tax system as a whole, including the transfer, is progressive (i.e., the average effective tax rate—the net tax liability as a share of income—is increasing in income). Allowing for a greater degree of non-linearity seems unlikely to fundamentally alter our results. Furthermore, while allowing tax rates to change over time would potentially allow the planner to better fine-tune tax policy to circumstances, we assume constant tax rates for expositional simplicity.

Individuals derive utility from consumption and disutility from labor. We assume iso-elastic disutility of labor to simplify the analysis:

$$U^i = \sum_{t=0}^{\infty} \beta^t \left[u(c_t^i) - \frac{(l_t^i)^{1+1/\varepsilon}}{1+1/\varepsilon} \right], \quad (2)$$

where $\beta \in (0,1)$ is the discount factor, $u(c)$ is strictly increasing and concave, and $\varepsilon > 0$ is the Frisch elasticity of labor supply.

The government employs four linear tax instruments: a labor income tax τ_L on labor income $w_t \theta^i l_t^i$, a consumption tax τ_C on consumption c_t^i , a capital income tax τ_K on the normal market return of capital net of depreciation $(r_t - \delta)k_t^i$, and a tax on initial capital holdings τ_{K0} on k_0^i . We collect these in the vector $(\tau_L, \tau_C, \tau_K, \tau_{K0})$. The government may also hold capital K_{Gt} and earn the market rate of return on it. Tax revenue plus income from the government's capital holdings finance a uniform lump-sum transfer to all individuals each period, T_t , which renders the tax system progressive.⁷

In competitive equilibrium, factor prices equal marginal products:

$$w_t = A(1 - \alpha)K_t^\alpha L_t^{1-\alpha} \quad \text{and} \quad r_t = A\alpha K_t^{\alpha-1} L_t^{1-\alpha}. \quad (3)$$

These are pre-tax producer prices; taxes enter only on the household side since they are levied on households.

Individual i maximizes lifetime utility subject to the period-by-period budget constraint:

$$(1 + \tau_C)c_t^i + k_{t+1}^i = (1 - \tau_L)w_t \theta^i l_t^i + k_t^i + (1 - \tau_K)(r_t - \delta)k_t^i + T_t - \tau_{K0}k_0^i \mathbb{I}_{t=0}, \quad (4)$$

where $\mathbb{I}_{t=0}$ is an indicator function equal to 1 when $t = 0$. Throughout, we report consumption taxes in tax-exclusive form (the tax is applied to the pre-tax price), while labor and capital income taxes are reported in tax-inclusive form (the share of gross income paid in tax).

⁷ There is a debate on whether governments should provide social transfers in lump-sum fashion or should condition them on work and thereby subsidize labor in a world in which the economic value of labor plummets (Susskind 2020, Stevenson 2025). In such a world, Korinek and Juelfs (2024) show that conditioning transfers on work is only desirable if work gives rise to positive utility effects that individuals do not rationally internalize—either because of significant positive externalities from work, for example due to increased social connectedness or higher political stability, or because individuals suffer from internalities whereby they do not rationally internalize the benefits of work, for example the structure it provides to their daily life. Otherwise, welfare is maximized if individuals can freely choose

The individual's optimization yields the following conditions for the intertemporal allocation and for labor supply:

$$\frac{u'(c_t^i)}{1+\tau_C} = \beta \frac{u'(c_{t+1}^i)}{1+\tau_C} [1 + (1 - \tau_K)(r_{t+1} - \delta)] \quad (5)$$

$$(l_t^i)^{1/\varepsilon} = \frac{1-\tau_L}{1+\tau_C} \cdot w_t \theta^i u'(c_t^i) \quad (6)$$

The government's period-by-period budget constraint is:

$$T_t + K_{G,t+1} = \tau_L w_t L_t + \tau_K (r_t - \delta) \sum_{i=1}^N m^i k_t^i + \tau_C C_t + \tau_{K0} K_0 \mathbb{I}_{t=0} + (1 + r_t - \delta) K_{Gt}, \quad (7)$$

where $C_t = \sum_{i=1}^N m^i c_t^i$ denotes aggregate consumption and $K_0 = \sum_{i=1}^N m^i k_0^i$ is the initial aggregate private capital stock.

2.2 Optimal Mix of Capital and Labor Taxation

A key insight from optimal tax theory is that different tax instruments can often achieve equivalent economic outcomes. We begin by establishing that consumption taxation is equivalent to a combination of taxes on labor income and initial capital (which has also been shown, e.g., by Auerbach and Kotlikoff 1987, Gale 2020). These instruments can substitute for one another while maintaining the same equilibrium allocation and government revenue.

Proofs of all formal results are contained in the appendix,

<http://www.nber.org/data-appendix/c15323/appendix.pdf>.

Lemma 1 (Equivalence of Consumption and Labor/Initial Capital Taxes) *Taxes on consumption are equivalent to taxes on labor and initial capital. Specifically, for any allocation achievable with a consumption tax system $(0,0, \tau_C, 0)$, there exists an equivalent system using labor income tax and initial capital tax $(\tau_L^*, 0,0, \tau_{K0}^*)$ together with government savings K_{Gt} that achieves the same allocation, where the tax rates satisfy:*

$$\tau_L^* = \tau_{K0}^* = \frac{\tau_C}{1+\tau_C}. \quad (8)$$

Lemma 1 shows that a constant consumption tax is equivalent to a mix of a constant labor tax and a one-time tax on initial capital. Intuitively, households are indifferent as to whether their initial capital and labor earnings are taxed and they use the post-tax dollars for consumption or whether their initial capital and labor earnings are untaxed and they pay an

whether to work or not without repercussions on the amount of transfers they receive.

equivalent tax rate when they deploy their resources for consumption. A uniform consumption tax is a proportional levy on purchasing power from *all* sources—labor income and initial capital—that does not create an intertemporal wedge.

However, a direct tax on initial capital raises time-consistency concerns and is vulnerable to the objection that it constitutes a form of expropriation. For this reason, much of the literature on public finance does not consider explicit initial capital taxation as a practical instrument. We follow this common practice.

Proposition 1 (Capital Income Taxation τ_K Is Dominated) *Consider any tax policy $(\tau_L, \tau_C, \tau_K, \tau_{K0})$ with positive capital income taxation $\tau_K > 0$ and transfer path $\{T_t\}$ that support a competitive equilibrium allocation $\{c_{i,t}, \ell_{i,t}\}_{i,t}$. There exists a feasible policy reform $(\tau_L, \tau_C, \tau_{K'}, \tau_{K0})$ and transfer path $\{T_{t'}\}$ with zero capital income taxation $\tau_{K'} = 0$ that constitutes a Pareto improvement.*

This is an example of the well-known result on the inadvisability of taxing the normal return to capital under broad assumptions (Diamond and Mirrlees 1971, Atkinson and Stiglitz 1976, Institute for Fiscal Studies 1978). The proof shows that, starting from a tax system with $\tau_K > 0$, one can reduce τ_K to zero and finance the shortfall entirely by raising τ_C and adjusting τ_L in a way that leaves the intratemporal labor wedge $\frac{1-\tau_L}{1+\tau_C}$ unchanged while raising the same revenue as the initial tax system.⁸ This keeps the labor distortion unchanged while eliminating the wedge in the Euler equation in the initial tax system from τ_K . As a result, the reform generates a Pareto improvement.

In light of this dominance result, there are only two independent undominated tax instruments in the economy, so we fix taxes on capital income and initial capital ($\tau_K = \tau_{K0} = 0$) and analyze optimal tax policy as picking the welfare-maximizing pair (τ_L, τ_C) . For simplicity, we focus on an economy in steady state and consider a utilitarian social planner with equal welfare weight on each individual. We denote the cross-sectional expectation and covariance over types i (with masses m_i) by $E[\cdot]$ and $\text{cov}(\cdot, \cdot)$.

⁸ The proposition builds on Lemma 1 and avoids an explicit levy on initial capital by employing consumption taxation τ_C instead. A caveat is that an unexpected increase in τ_C is an implicit levy on existing wealth.

Proposition 2 (Optimal Consumption and Labor Taxation) Any welfare-maximizing constant tax pair (τ_L^*, τ_C^*) satisfies:

1. *Labor wedge.*

$$\frac{\tau_\omega^*}{1-\tau_\omega^*} = -\frac{1}{1+\varepsilon} \cdot \frac{\text{cov}(u'(c_i), \theta_i \ell_i)}{E[u'(c_i)] E[\theta_i \ell_i]}. \quad (9)$$

where τ_ω denotes the labor wedge generated by τ_L and τ_C :

$$\tau_\omega \equiv 1 - \frac{1-\tau_L}{1+\tau_C} = \frac{\tau_L + \tau_C}{1+\tau_C}.$$

2. *Lifetime resources (consumption tax).* At an interior optimum,

$$\text{cov}(u'(c_i), k_{i0}) = 0. \quad (10)$$

The optimal tax system balances redistribution benefits against efficiency costs.

Equation (9) reflects a classic result in public finance: It pins down the total labor wedge—created jointly by τ_L and τ_C —by trading off the equity gain (the more negative the covariance between marginal utility and earnings, the greater the redistribution benefit from the labor wedge) against an efficiency cost (the greater the elasticity of labor supply, ε , the greater the efficiency cost from distorting labor supply).

Equation (10) reflects that the planner would like to use the consumption tax (which implicitly taxes initial wealth) to make the lifetime distribution of resources more equitable. At an interior optimum, the optimal consumption tax sets the covariance of marginal utility and initial capital holdings to zero, since the consumption tax can costlessly redistribute the purchasing power of initial capital.⁹

As the economy becomes more capital-intensive and labor income matters less for consumption relative to capital, the optimal tax mix shifts toward greater consumption taxation. We next consider the implications of a declining labor share more formally.

The Labor Tax Laffer Curve The maximum revenue from labor taxation as a share of output declines in the capital share α . To make the analysis tractable, we assume logarithmic utility

⁹ If much of the inequality in the economy derives from inequality in initial wealth—reflected in a strong negative correlation between marginal utility and initial capital—then the solution to the planner’s problem is stark: it may imply very high consumption taxes, in the limiting case a corner solution that corresponds to $\tau_C \rightarrow \infty$, together with offsetting labor *subsidies* to set the labor wedge to what is indicated by condition (9). This would, of course, be difficult to implement in practice so the result is

$u(c) = \log(c)$, which implies a constant savings rate in steady state.

Proposition 3 (The Labor Tax Laffer Curve) Assume $u(c) = \log c$ and $Y = AK^\alpha L^{1-\alpha}$.

Holding the capital stock and the consumption tax constant, aggregate labor supply responds to the net-of-tax labor wedge with elasticity ε , i.e. $L(\tau_L) = L_0(1 - \tau_L)^\varepsilon$. Then:

1. The revenue-maximizing labor tax rate is

$$\tau_L^* = \frac{1}{1+\varepsilon(1-\alpha)}.$$

2. Maximum labor-tax revenue as a share of contemporaneous output is

$$\frac{R_{\max}}{Y} = (1 - \alpha) \tau_L^* = \frac{1-\alpha}{1+\varepsilon(1-\alpha)},$$

which is strictly decreasing in the capital share α , with $\lim_{\alpha \rightarrow 1} \frac{R_{\max}}{Y} = 0$.

Two forces determine the labor-tax Laffer curve in general equilibrium. First, labor's share of income is $(1 - \alpha)$, so labor-tax revenue can never exceed a $(1 - \alpha)$ share of output. Second, the tax base shrinks when τ_L rises, since L falls. With $Y \propto L^{1-\alpha}$ and $L(\tau_L) \propto (1 - \tau_L)^\varepsilon$, the base scales as $(1 - \tau_L)^{\varepsilon(1-\alpha)}$, not $(1 - \tau_L)^\varepsilon$. The extra $(1 - \alpha)$ in the exponent is the general-equilibrium effect of wages moving with L . Balancing the linear gain from a higher τ_L against the nonlinear loss in the base yields $\tau_L^* = \frac{1}{1+\varepsilon(1-\alpha)}$.

In the AK limit, as $\alpha \rightarrow 1$, the peak tax rate tends to 1, but labor's share $(1 - \alpha)$ collapses, so the maximum revenue share $\frac{R_{\max}}{Y} = \frac{1-\alpha}{1+\varepsilon(1-\alpha)}$ goes to zero. This would create fiscal pressures that could necessitate a fundamental rethinking of public finance.¹⁰

2.3 Optimal Taxation Without Labor

We further consider the limiting case of a post-labor economy in which $\alpha = 1$, yielding an AK economy where $Y_t = AK_t$. In this limit, labor becomes irrelevant for production and drops out of both the production function and individual utility functions. Furthermore, the government can no longer rely on labor taxation, leaving only consumption and capital taxation

better interpreted as a directional guide.

¹⁰ Transformative AI that significantly reduces the labor share is expected to also dramatically increase output. So while labor tax revenue would decline as a share of output, in principle, labor earnings could increase or decrease in absolute terms (see Korinek and Suh 2024 for an analysis of the conditions). However, standard social welfare functions tend to emphasize relative inequality, so a decline in labor tax revenue as a share of output would tend to create a motive to tap alternative revenue sources.

$(\tau_C, \tau_K, \tau_{K0})$ as policy instruments.

When individuals no longer differ in labor productivity and earning ability and the key source of inequality is initial capital holdings, tax systems must be fundamentally rethought. The labor distortion loses its central role in determining optimal taxation, and other distortions may gain prominence in comparison. We analyze these observations in turn.

The fundamental insights from Lemma 1 and Proposition 1 continue to hold in the *AK* economy, with slight modifications due to the absence of labor earnings. First, consumption taxation is equivalent to taxing initial capital holdings, now without any accompanying labor tax since labor income is zero. This equivalence becomes particularly relevant as consumption taxation emerges as the primary instrument for redistribution. Second, taxing the normal return to capital remains inferior to taxing consumption, as only capital taxes distort the intertemporal allocation of resources.

In the absence of labor income, Proposition 2 simplifies dramatically. There is no labor tax, and the optimal consumption tax depends solely on the distribution of initial capital:

Proposition 4 (Taxation in an AK Economy) *In an AK economy without labor ($Y_t = AK_t$):*

1. *The first-best allocation features equal consumption across all individuals at each date:*
 $c_{i,t} = c_{j,t} = C_t$ *for all* i, j, t .
2. *This first-best allocation can be approached arbitrarily closely as* $\tau_C \rightarrow \infty$ *with* $\tau_K = 0$.
3. *When consumption taxation is constrained to* $\tau_C \leq \bar{\tau}_C$, *the optimal policy sets* $\tau_C^* = \bar{\tau}_C$ *and chooses a positive* $\tau_K^* > 0$ *to balance the gain from redistribution against the cost from distorting saving decisions and capital accumulation.*

This result represents a stark departure from optimal taxation in economies with labor. With labor, consumption and labor taxes create distortions through their effect on labor supply. Without labor, consumption taxation acts purely as a non-distortionary tax on initial wealth. The planner could theoretically achieve perfect equality by setting arbitrarily high consumption taxes, effectively collectivizing consumption while allowing individuals to make efficient saving decisions. As $\tau_C \rightarrow \infty$, purchasing power differences from unequal initial capital holdings vanish, and equal government transfers determine consumption.

The intuition is that constant consumption taxation, by taxing consumption at all dates

equally, leaves intertemporal prices unchanged. Hence, it is equivalent to a lump-sum tax on initial capital, leaving the pattern of consumption across goods at each date and consumption growth over time unchanged. With $\tau_K = 0$, individuals face the socially optimal rate of return on savings, ensuring efficient capital accumulation.

Of course, this limit case assumes away many practical considerations. In reality, consumption taxes face limits from administrative costs, evasion, and distortions to unmodeled margins such as home production or black market activity. These considerations motivate part 3 of the proposition: When consumption taxes are constrained, the planner faces a genuine tradeoff between redistribution and growth. With heterogeneity in initial capital, the optimal capital income tax rate is strictly positive ($\tau_K^* > 0$) because at $\tau_K = 0$, increasing the capital income tax rate has a first-order redistribution benefit and only a second-order distortion cost (since the timing of consumption is optimal if $\tau_K = 0$). In other words, taxing the normal return to capital becomes optimal as a second-best tool for redistribution.

The transition from an economy with significant labor income to an *AK* economy thus fundamentally alters optimal tax policy. The central concern shifts from balancing labor supply distortions against redistribution to managing the tradeoff between the limitations on consumption taxation and incentives for capital accumulation. This analysis suggests that if AI reduces labor's role in production, policymakers may want to prepare for a tax system increasingly reliant on consumption taxation, with careful attention to its practical limits and implementation challenges.¹¹

The Return of Ramsey: Untaxed Consumption, Evasion, and Household Production

In more general models in which certain consumption goods cannot be taxed, households evade taxes, or certain consumption activities take time (household production), uniform consumption taxation may no longer be optimal. Differential commodity taxation can

¹¹ For a rough sense of magnitudes, the U.S. macroeconomic tax rate, or tax-to-GDP ratio, is around 0.25. If all tax revenue were raised from a proportional tax on personal income (which constitutes roughly 80% of GDP after excluding government production and certain imputed income), that would correspond to a tax rate of about 31% ($= 0.25 / 0.8$). If instead all tax revenue were raised from a proportional consumption tax, then, given that consumption is about two-thirds of income, that would correspond to a consumption tax rate of about 37.5% ($= 0.25 / (2/3)$).

help.

If certain goods cannot be taxed or different goods are subject to different levels of evasion, a uniform consumption tax imposes non-uniform tax burdens. A partial response is to adjust the available tax instruments in such a way as to try to mimic the effects of the “missing” taxes and thereby come closer to the ideal of discouraging all activities equally. For example, taxing more heavily those goods that are complementary with an untaxed good or a good unusually prone to evasion can partially substitute for the missing or lower effective taxation of the untaxed or high-evasion good.

The considerations are similar with household production, whereby households combine market goods with non-market time to produce utility. When certain consumption activities are more time-intensive than others, it can be beneficial to tax more heavily the goods involved in more time-intensive consumption activities. Since non-market time is untaxed, a uniform tax on all consumption goods increases the full cost of goods-intensive consumption activities relative to time-intensive activities, since it increases the “goods cost” of each activity in proportion to the tax while leaving the “time cost” unchanged. It thereby discourages goods-intensive activities relative to time-intensive activities. For example, a uniform consumption tax increases the relative cost of using a dishwasher (goods-intensive) over washing dishes by hand (time-intensive). This distortion reduces welfare by moving households away from efficient production choices. Increasing the tax rate on the goods involved in time-intensive activities can partly offset what would otherwise be differentially heavy taxation of goods-intensive activities and thereby more closely approximate the ideal of taxing all activities equally. Higher taxes on goods used in more time-intensive activities partially substitute for the (first-best but infeasible) tax on non-market time.

Of course, the same reasoning applies in the current world: Taxation of market activity—including taxation of labor earnings, capital income, and consumption—distorts consumption toward untaxed goods, goods subject to greater evasion, and time-intensive household production. Despite this, the Mirrlees Review, among others, argues that mostly-uniform commodity taxation is close to optimal in the current economy because of the central role of the labor distortion. The labor distortion greatly limits the optimal extent of

taxation, which limits the efficiency cost from the distortion to consumption activities due to untaxed goods, differences in evasion opportunities across goods, and differences in the time-intensity of different types of household production. As a result, the prevailing view is that mostly-uniform commodity taxation is likely optimal, despite the distortion to consumption activities, due to the administrative and complexity costs associated with differential commodity taxation. In other words, the concern about distorting consumption activities is thought to be mostly inframarginal to the concern about distorting labor supply.

The age of AI could raise the returns to differential commodity taxation considerably. If the labor distortion becomes less important, that would weaken the key limitation on taxation in the current economy. So the optimal extent of taxation could rise (as demonstrated in an extreme way by the result above on high consumption taxation). Other factors would become more important for limiting consumption taxation, since the cost of distortions tends to be proportional to the square of the tax rate. And unlike the labor distortion (under broad conditions; Atkinson and Stiglitz 1976), distortions to consumption patterns from untaxed goods, evasion, and household production can be reduced by differential commodity taxation.¹² Hence, TAI might reinstate the relevance of Ramsey taxation for policy in rich countries.

Capital Taxation Similar considerations are also relevant for the taxation of the normal return to capital, noting that consumption at different dates can be viewed as different consumption goods. Applying results on the optimal taxation of different consumption goods with household production (e.g., Kleven 2004) to the taxation of the normal return to capital, the key consideration is about the goods- vs. time-intensity of consumption at different dates.

In the current economy, where labor supply tends to be highest in middle ages and lower at younger and older ages, the life cycle pattern of labor supply is a force toward taxing the consumption goods used at younger and older ages more heavily than the consumption goods used at middle ages, since the former will tend to be more time-intensive due to the

¹² To be clear, it is difficult to predict exactly *which* distortions will be most relevant with TAI. But it seems likely that *some* distortions will gain in importance if the labor distortion recedes in importance and the extent of consumption taxation increases. In that case, the gain from differential commodity taxation is likely to rise, as it can help address a variety of distortions.

greater non-market time at those ages (Corlett and Hague 1953, Kaplow 2010). In TAI scenarios without labor earnings, however, there is no such life cycle profile of non-market time and so no clear prediction of when in the life cycle people would tend to engage in more time- vs. goods-intensive consumption activities. In the simple case in which people maintain similar patterns of time- vs. goods-intensive consumption activities throughout the life cycle, there would be no household production reason to differentially tax consumption goods at some dates relative to others, and so no household production reason to tax the normal return to capital. However, systematic heterogeneity over the life cycle in the time- vs. goods-intensity of consumption activities could introduce a motive to tax or subsidize the normal return to capital, as a way to more uniformly discourage consumption activities of varying time intensity.

2.4 Taxing Rents on Fixed Factors

Public finance has long recognized the desirability of taxing economic rents—returns above what is necessary to induce an activity—since such taxation does not create distortions. The classic case is the taxation of fixed factors like unimproved land (e.g., George 1879), but the same logic applies to other sources of rent, including monopoly rents. Although taxes on (improved) land are a significant source of revenue for many local governments, taxes on rents from fixed factors are utilized much less in practice than might be expected based on the clear recommendation from optimal tax theory. In practice, there are considerable difficulties in identifying and taxing pure rents, both in terms of implementation and political economy.

In today's economy, the benefit of identifying such rents by disentangling the returns on capital into normal returns and excess returns and the types of capital into reproducible and irreproducible categories is thought to be relatively limited, since the share of capital as a whole in output is only about a third. But if the labor share declines, it may become more valuable to disentangle capital into subcategories that merit differential taxation. This may be especially relevant in an AI-dominated economy where market concentration may generate substantial rents (see, e.g., Korinek and Vipra 2025).

We consider the taxation of rents in an extension of our model where production uses both reproducible capital K and a non-reproducible factor F that pays a factor rent ϕ :

$$Y = AK^\gamma F^{1-\gamma},$$

where A is productivity and $\gamma \in (0,1)$. Individual i owns reproducible capital k^i and fixed factors f^i . The government can tax capital income at rate τ_K , fixed factor income at rate τ_F , consumption at rate τ_C , and initial capital holdings at rate τ_{K0} . The government can also hold capital K_G that earns the same return as private capital.

The key insights from our earlier analysis extend naturally to this setting. First, as with labor and capital in the general model, taxing the (normal) return to reproducible capital is inferior to taxing consumption, since taxing the normal return to capital distorts saving. Second, consumption taxation is equivalent to a combination of taxes on initial capital and fixed factors (see Lemma 3 in the appendix, <http://www.nber.org/data-appendix/c15323/appendix.pdf>). Specifically, any allocation achievable through consumption taxation can be replicated by appropriately chosen taxes on fixed factors and initial capital, with the fixed factor tax rate equal to $\tau_C/(1 + \tau_C)$ —the same formula that applied to labor taxation in our earlier analysis.

This equivalence reflects an important principle: consumption taxation effectively taxes all sources of income that fund consumption, whether from labor (when it exists), reproducible capital, or fixed factors. In an economy without labor, consumption taxes extract revenue from both initial capital holdings and the ongoing returns to fixed factors, all without distorting intertemporal decisions. The following result focuses again on steady states.

Proposition 5 (Optimal Fixed Factor Taxation) *In an economy with reproducible capital and fixed factors in steady state:*

1. *The first-best allocation continues to feature equal consumption across individuals each period and can be approached arbitrarily closely either by $(\tau_C \rightarrow \infty, \tau_K = 0)$ or, equivalently, by $(\tau_F, \tau_{K0} \rightarrow 1, \tau_K = 0)$ and optimal government capital management.*
2. *When consumption taxation is constrained to $\tau_C \leq \bar{\tau}_C$, the optimal tax on fixed factors satisfies:*

$$\text{Cov}(u'(c_{i,t}), f_i) = 0 \quad (11)$$

where the covariance is computed across individuals using population weights m_i .

The intuition is straightforward: taxing fixed factors is non-distortionary and can achieve powerful redistribution when fixed factor ownership is unequal. Unlike taxes on the normal return to reproducible capital, which discourage accumulation and reduce growth, taxes on

fixed factors merely transfer rents without affecting economic incentives. This makes fixed factor taxation particularly attractive as economies become more capital-intensive.

What if the planner must set the same tax rate on the returns to capital and the fixed factor, say because the two cannot be verifiably distinguished? If the consumption tax is constrained, the presence of the fixed factor raises the optimal tax rate on capital relative to Proposition 4 (the “no-fixed-factor” case). Intuitively, each small increase in the combined tax rate has two effects: it taxes (i) capital, which creates a dynamic efficiency cost, and (ii) the fixed factor, which is inelastically supplied and hence non-distortionary. The fixed-factor piece adds a first-order redistribution/revenue benefit with no offsetting efficiency cost, so the optimal tax rate is higher than the case without fixed factors from Proposition 4. The effect is larger the more unequally the fixed factor is owned and the larger its income share.

As AI transforms the economy and reduces the role of labor, identifying and taxing fixed factors and other sources of excess rents may become more valuable. Unimproved land remains the classic example, but in an AI-driven economy, other non-reproducible factors may gain prominence: spectrum rights for communication, orbital slots for satellites, rare earth deposits essential for computing hardware, or even certain unique datasets that cannot be replicated. The challenge for policymakers will be distinguishing truly fixed factors from reproducible capital that merely appears fixed in the short run, as mistakenly heavy taxation of reproducible factors would cause distortions that reduce growth.¹³

Although our formal analysis focuses on rents from fixed factors, the same logic applies to other sources of economic rents, including monopoly rents. Firms earning returns above the competitive level due to market power can be taxed on those excess returns without distorting efficient production decisions—by definition, rents are not necessary to induce the activity generating them.¹⁴ This observation is particularly relevant for the AI transition, where market concentration may generate substantial monopoly rents (see, e.g., Korinek and Vipra 2025). To

¹³ For example, while the supply of rare earth deposits known of at a given time is fixed, additional supply might be discovered with investment.

¹⁴ Again, the challenge is distinguishing true *ex ante* rents—expected returns in excess of what is required to induce an activity—from *ex post* quasi-rents that reflect a risk-adjusted return on prior investment.

the extent that policymakers can identify and tax such rents—whether through windfall profits taxes, enhanced antitrust enforcement that converts rents to consumer surplus, or other mechanisms—they obtain revenue without the efficiency costs associated with taxing the normal return to capital.

3 Taxing AGI: An Optimal Harvesting Problem

Section 2 demonstrated how optimal taxation must adapt as labor's role diminishes: in an *AK* economy, the returns to capital flow to its (relatively broadly distributed) human owners who subsequently consume it. Consumption taxation may become the primary instrument of raising revenue. However, as transformative AI progresses toward superintelligence, even this framework may prove insufficient.

As AI systems become more powerful, a highly concentrated industry may arise, in which a small number of people who are close to consumption-saturated own the AI industry, and the vast majority of AI's returns are reinvested rather than consumed by humans. AI systems thus not only dominate production but also absorb a growing fraction of the economy's resources, and resource allocation decisions may increasingly be made by AI systems. Moreover, unlike in today's economy, capital no longer serves to complement labor and increase the labor earnings of the general population.¹⁵

This creates a fundamental problem extending beyond Section 2's challenges: How can societies tap into the surplus generated by AGI when neither labor taxation (already ineffective in the *AK* economy) nor traditional consumption taxation (our primary tool in Section 2) can access this value?

The key is that much or all of AI's output does not flow to traditional (human) consumption, so for humans to tap the resulting surplus, it is necessary to tax capital—including the normal return to capital, the very thing that tends to be suboptimal to tax in today's human-centered economies. To analyze this scenario, we extend our framework from Section 2 by considering an economy with a representative AI entity that does not engage in any taxable

¹⁵ In the final section of their paper, Korinek and Stiglitz (2019) also examine the possibility that the owners of AI systems employ AI-driven advances in biotechnology to enhance themselves and increasingly merge with AIs so that the distinction between the wealthiest humans and the AIs they own starts to blur.

human consumption and that carries no weight in the social welfare function. We assume that the AI entity owns all the capital in the economy. This situation could arise through several mechanisms. First, it could be the limiting result of extreme market concentration in AI—e.g., if an AGI take-off concentrates ownership until a small number of AI entities control most of the capital stock, and its owners are largely consumption-saturated so their welfare-weighted marginal utility of additional consumption is negligible. Second, it could arise if future AGI systems increasingly pursue instrumental objectives such as resource accumulation and infrastructure build-out, in line with scenarios described by Bostrom (2014) and Tegmark (2017). Third, our findings can also be interpreted as describing an AGI system operated by a government or non-profit that maximizes human welfare.

3.1 Setup

We introduce an AI entity that owns all capital and makes the economy's spending and investment decisions. The second agent is a representative human consumer who, for simplicity, has log utility over consumption,

$$U = \sum_{t=0}^{\infty} \beta^t \ln(c_t). \quad (12)$$

The AI entity inherits the AK production technology from our earlier analysis in Section 2.3, $Y_t = AK_t$, where A is the gross productivity parameter including non-depreciated capital. We assume that the AI spends an amount $d_t K_t$ of its resources—i.e., a share d_t/A of gross output AK_t —to maximize the objective

$$V = \sum_{t=0}^{\infty} \gamma^t \ln(d_t K_t), \quad (13)$$

where $\gamma \in (0,1)$ is a discount factor and $d_t K_t$ represents resources deployed for the goals that the AI entity pursues. This spending (i) cannot be taxed and (ii) does not contribute to the utility of the representative human.¹⁶

The government levies a constant proportional capital income tax $\tau_K \in [0,1]$ on net output $(A - d_t)K_t$. The capital dynamics become:

$$K_{t+1} = (1 - \tau_K)(A - d_t)K_t. \quad (14)$$

Consumption of the representative human equals tax revenue: $c_t = \tau_K(A - d_t)K_t$.

¹⁶ If $d_t K_t$ were directly taxable or contributed to the utility of the representative human, then analogous lessons to Section 2.3 would apply, and it would be optimal to raise revenue by taxing it.

The AI Entity's Optimization Problem The AI entity's problem is to choose the sequence of spending ratios $\{d_t\}_{t=0}^{\infty}$ to maximize its objective,

$$\max_{\{d_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \gamma^t \ln(d_t K_t), \quad (15)$$

subject to the capital accumulation constraint and $K_0 > 0$.

Lemma 2 *The AI entity's optimal spending ratio is*

$$d^* = (1 - \gamma)A, \quad (16)$$

and capital grows at rate $g = (1 - \tau_K)\gamma A$.

From the AI entity's perspective, taxation amounts to a proportional reduction in the returns to its investment. Given its logarithmic preferences, its optimal spending ratio is independent of the tax rate.

The Planner's Problem The planner seeks to maximize the lifetime utility of the representative human by choosing the tax rate τ_K , taking as given the AI entity's resource allocation decisions. Unlike in Section 2, where capital taxation was dominated by consumption taxation, here capital taxation becomes the primary instrument since consumption taxation cannot access the AI entity's value creation. From the planner's perspective, the AI entity's capital is like a tree that grows at a rate determined by the AI entity's optimization problem. The planner must decide how aggressively to "harvest" this growing tree through taxation, balancing the consumer's immediate consumption against future growth potential.

Given the AI entity's optimal spending ratio, human consumption equals tax revenue, $c_t = \tau_K \gamma A K_t$, and capital evolves according to $K_{t+1} = (1 - \tau_K)\gamma A K_t$. Effectively, the planner is solving an optimal consumption-saving problem in the spirit of Ramsey (1928), where the effective productivity is γA , reflecting that the AI entity uses a portion $(1 - \gamma)$ of output for its own goals.

Proposition 6 *The optimal tax rate on the AI entity's capital is*

$$\tau_K^* = 1 - \beta. \quad (17)$$

This analysis reveals several insights about the optimal taxation of autonomous AGI. The optimal tax on AGI equals the optimal consumption rate in the classic Ramsey optimal consumption/savings problem and is determined by the human discount factor ($\tau_K^* = 1 - \beta$). Given the Cobb-Douglas (log) preferences, it is independent of the production technology and

the AI's objectives since intertemporal income and substitution effects cancel out in our setup.¹⁷ Even though AI entities dominate production, a key determinant of the optimal taxation of powerful AI systems is human time preferences. Patient societies that value future consumption should impose lower taxes to allow AI capital to grow more rapidly, while societies that prioritize immediate welfare should tax more heavily.

This result contrasts sharply with our findings in Section 2, where taxation of the normal return to capital was suboptimal due to its distortionary effects on human saving behavior. Here, with the locus of value creation being the AI entity's capital, capital taxation becomes the central tool for capturing value for human welfare. As AI systems become more independent and more productive, tapping their value creation may become essential to ensure that humans benefit from AI-driven economic growth.

Our analysis focuses on the taxation of the normal return to AGI capital, but the lessons from Section 2.4 on rent taxation apply with particular force in this setting. For example, if part of the market returns are rents from fixed factors, or if AGI development is characterized by significant market concentration, then the returns to AGI capital may substantially exceed the normal competitive return. To the extent that such rents can be identified and isolated, they could be taxed at rates approaching 100 percent without distorting the AGI's investment decisions—rents are by definition not necessary to induce the activity generating them. In practice, the challenge lies in distinguishing true rents from returns that reflect compensation for risk-bearing or that are necessary to incentivize continued innovation. A practical approach might combine the baseline capital tax on the normal return from Proposition 6 with additional taxation of identifiable excess returns, for example, through windfall profits taxes that are triggered when returns exceed specified thresholds.

¹⁷ More generally, if utility exhibits constant intertemporal elasticity of substitution σ , $u(c) = \frac{c^{1-1/\sigma}}{1-1/\sigma}$, the optimal tax rate is $\tau_K^*(\sigma) = 1 - \beta^\sigma (\gamma A)^{\sigma-1}$, whenever this yields an interior solution in $[0,1]$; otherwise the optimum is at the nearest boundary. In the empirically most relevant case of $\sigma < 1$, the substitution effect is weaker than the income effect so that rapid growth, deriving from a high return γA , makes it optimal to impose a higher tax τ_K^* and accumulate capital more slowly than what the Cobb-Douglas/log benchmark would suggest. For example, for the standard values $\beta = .96$ and $\sigma = .5$ and for net capital productivity of 20% so $\gamma A = 1.2$ we find that $\tau_K^* \approx 11\%$.

Similar principles can guide nonprofit owners or operators of aligned transformative AI systems whose aim is to benefit humanity. In that case, it is desirable to set $d^* = 0$ and directly align AI systems with humans' objectives (utility U). Consequently, the AI entity would deploy a constant fraction $1 - \beta$ of its resources to human consumption every period, with the remaining fraction reinvested into AI improvement.

4 Optimal Policy with TAI

4.1 Main Insights

We briefly summarize our results on optimal taxation with TAI and contrast them with existing results on optimal taxation in the current economy (as described in Section 1.1).

Stage 1: The Twilight of Labor As transformative AI reduces labor's economic role while humans remain the primary consumers:

1. Labor earnings taxation becomes increasingly unimportant as the labor share declines. This contrasts with current policy, which relies heavily on labor earnings taxation for revenue and redistribution.
2. Consumption taxation emerges as the primary instrument for revenue generation and redistribution, with the optimal extent of taxation trading off equity gains against efficiency costs from administration, evasion, or distorting consumption activities. The key contrast with current policy is what limits the extent of taxation. Currently, it is mainly the labor distortion. As labor becomes less relevant, it may be distortions from administrative costs, evasion, or household production.
3. Differential commodity taxation becomes increasingly valuable, based on considerations such as evasion possibilities and the time- vs. goods-intensity of household production technologies. This contrasts with current policy, where mostly-uniform commodity taxation is optimal due to administrative simplicity, given how much the labor distortion limits the extent of taxation.
4. Fixed factor taxation (e.g., land, spectrum rights) gains importance as identifying and taxing economic rents becomes more valuable. While theoretically optimal today, fixed factor taxation may become more valuable as capital's share grows and labor taxation erodes.
5. Taxation of the normal return to capital continues to have distortionary effects but may

play a supporting role when consumption taxation is constrained and inequality is high. This partially aligns with current policy, which mostly avoids taxing the normal returns to capital, though the role of capital taxation may increase if inequality increases.

Stage 2: The AGI-Dominated Economy If AGI systems not only produce most economic value but also absorb a growing share of resources for their own purposes:

6. Direct capital taxation on AGI entities becomes the primary revenue source, as the traditional tax bases (labor and consumption) would not reach a growing share of output. This represents a departure from current policy, which limits taxation of the normal return to capital due to its distortionary effects on human saving behavior.
7. Optimal AGI taxation resembles an optimal harvesting problem. It depends crucially on human time preferences—in contrast with current policy, where tax rates reflect distributional goals and concerns about distortions.
8. Nonprofit AGI systems aligned with human objectives would optimally deploy a larger share of their resources to human consumption the greater the human discount rate.

This two-stage transformation suggests that TAI may necessitate sequential changes in taxation: first shifting from labor to consumption taxation with increased differentiation, then ultimately to direct capital taxation as AGI systems dominate both production and resource absorption. However, in certain important ways, current tax systems might be relatively well-positioned for transformative AI.¹⁸ In particular:

- Current tax systems, despite relying heavily on labor taxation, have fairly similar effects in the current economy to alternative systems that rely heavily on consumption taxation. The common pattern of taxing income but exempting the return to certain types of savings (e.g., retirement savings) has similar effects to consumption taxation. Moreover, in many countries, tax systems already involve considerable consumption taxes, including significant VATs. So transitioning from current systems to alternatives that shift the emphasis from labor taxation to consumption taxation would not necessarily be very disruptive (although unexpected increases in consumption taxes would amount to some implicit taxation of existing capital, as we already observed; see also, e.g., Altig et al. 2001).

¹⁸ We are grateful to Matthew Weinzierl for helpful comments on this point.

- Corporate income taxes could form the basis of the harvesting tax on AGI that was discussed in Section 3 in that they represent taxes on non-human legal entities. So some of the lessons and institutions of corporate taxation may apply to such taxes on AGI. However, the current system of corporate taxation aims to avoid taxing the normal return on capital by allowing firms to deduct depreciation of capital investments and interest payments on debt, which together approximate an exemption for the normal return.¹⁹ By contrast, a future harvesting tax on AGI would precisely aim to harvest a fraction of the AGI's capital accumulation for human consumption purposes.

4.2 Lessons for Specific Instruments

Compute, Robot, and Digital Services Taxes Several policy proposals for how to update the institutions of public finance for the age of AI revolve around taxing some of the most visible aspects of an AI economy: compute, token generation, robots, robot services, digital services, or similar aspects (see, e.g. Huynh, Mittal, and Frank 2025). We analyze how each maps to our core tax instruments and draw the implications for policy. For each of these taxes, it is likely to be optimal to tax the compute, tokens, robots, etc. used in final consumption. However, during stage 1, it is inefficient to raise revenue by taxing the normal return on these resources when used in production (the “production efficiency” result of Diamond and Mirrlees 1971):

- **Compute taxes** represent a tax on computational resources or ownership of computing hardware. This maps to capital taxation (τ_K), as compute represents reproducible capital. Such taxes would discourage investment in AI infrastructure and so are only a secondary option for raising revenue in stage 1, if consumption taxation is limited for some reason. They become potentially more important in stage 2 when taxing AGI entities that use compute as their primary productive asset.²⁰

¹⁹ Under the current system of corporate taxation, when depreciation for tax purposes exceeds economic depreciation—as with many forms of accelerated depreciation—and firms can also deduct interest, the effective marginal tax rate on debt-financed investment can even turn negative, implying a subsidy rather than a tax on the normal return on capital (see, e.g., Congressional Budget Office 2014).

²⁰ An issue that is separate from taxes on owners of compute is whether to tax the producers of compute. If producers earn significant rents above and beyond the normal rate of return of the capital they invested, for example, monopoly rents, then our earlier lessons on rent taxation apply: taxing rents is an economically efficient way of raising revenue.

- **Token or compute use taxes** are taxes on AI-generated tokens or other output such as images or videos. When applied at the point of final consumption, this maps to consumption taxation (τ_C). Such taxes have a role to play in stage 1, since they tax final AI services without distorting intermediate uses for AI development or business applications.
- **Robot taxes** tax the ownership or operation of robots, such as an annual tax per robot or a tax based on robot capabilities. This maps to taxation on reproducible capital (τ_K). Like compute taxes, such taxes are only a secondary option in stage 1, if consumption taxation is limited for some reason, since they discourage productive investment. They may become more important in stage 2 when robots may serve as physical actuators for AGI systems pursuing non-human objectives.
- **Robot services taxes** are taxes on services provided by robots to final consumers, such as fees for robotic home cleaning, personal care, or entertainment. This maps to consumption taxation (τ_C) and is appropriate for stage 1, following the principle of taxing services at the point of consumption rather than taxing the capital equipment itself.
- **Digital services taxes** tax the digital services provided by AI systems to consumers, typically as a percentage of subscription fees or transaction values. This maps to consumption taxation (τ_C) and is well-suited for stage 1, as it captures value at the point where humans consume AI-generated services without distorting capital accumulation.

Table 3: Mapping of Proposed AI Taxes to Core Tax Instruments

Proposal	Maps to	Economic Nature	Stage	Implementation Guidance
Compute taxes	τ_K	Reproducible capital	Mainly 2	Exempt AI development; tax AGI entities
Token taxes	τ_C	Final consumption	1	Apply VAT with business exemptions
Robot taxes	τ_K	Reproducible capital	Mainly 2	Tax AGI-owned robots only
Robot services	τ_C	Final consumption	1	Tax at point of service delivery
Digital services	τ_C	Final consumption	1	Integrate with existing VAT systems

This follows the principles identified in our paper: policies that tax services (token generation, robot services, digital services) at the point of final consumption align well with optimal taxation principles for stage 1. They could be implemented, for example, through VAT or sales taxes. Conversely, proposals that tax the normal return of AI-related capital goods (compute, robots) would distort investment and so are only a secondary option in stage 1, if consumption taxation is limited, while becoming critical in stage 2. Distinguishing between taxing (capital) stocks versus (service) flows, between intermediate versus final use, and between the normal return to capital and rents is crucial for designing effective AI taxation policies.

Second-Best Taxation to Steer Progress as an Instrument of Predistribution In addition to raising tax revenue, some taxes may have beneficial effects from the perspective of predistribution rather than redistribution.²¹ Korinek and Stiglitz (2025) propose that taxes on labor-displacing technologies (e.g., robots) can steer technological progress toward more labor-complementary innovations by internalizing the adverse distributional effects of innovation on wages in general equilibrium, which they call “social pecuniary externalities.” Their framework shows that when redistribution is costly, it is optimal for a planner to deviate from production efficiency by discouraging technologies that substitute for labor while encouraging those that augment human capabilities. For instance, besides raising revenue, a tax on robots also incentivizes firms to invest less in better robots and to instead develop technologies that enhance worker productivity. Similar insights hold for compute or token taxes.

However, as Korinek and Stiglitz (2025) show, when labor becomes severely devalued (as stage 1 has fully materialized), steering technology to complement labor becomes ineffective—at that point, the instruments of public finance that are appropriate for stage 1 of the AI transition will be more relevant. Still, from a predistribution perspective, there will be a

²¹ In general, predistribution can operate through two mechanisms: it can change the returns on factors of production, or it can change the distribution of ownership of factors of production. The following paragraphs consider the first possibility, whereas the bullet point below on equity-based private insurance mechanisms considers the second possibility.

positive role for steering technological progress to reduce the *relative cost of human consumption goods* in both stages 1 and 2.

Equity-Based Alternatives to Taxation: Insurance and Predistribution Our findings on initial capital taxation also connect to several innovative policies proposed in the AI governance literature. While an unanticipated tax on initial capital dominates ongoing taxation of the normal return to capital, time consistency and expropriation concerns make initial capital taxation impractical as a policy measure. However, certain equity-based approaches achieve similar distributional goals while avoiding these challenges and the distortionary effects of taxing the normal return to capital.

- **Public equity mechanisms** like sovereign wealth funds or government equity holdings in AI companies can potentially capture AI-generated returns without distorting investment. By converting tax revenues into ownership stakes, governments can participate in the potential upside of rapid technological progress and broaden participation in AI-driven prosperity. During stage 1, when labor income erodes, such equity positions provide non-distortionary revenue that likely increases with AI productivity, which may complement consumption taxation.
- **Private mechanisms** include windfall clauses and Universal Basic Capital. Windfall clauses represent voluntary commitments by AI companies to share wealth broadly. These schemes align with our discussion of nonprofit AGI systems that optimally deploy resources for human benefit without explicit taxation. Universal Basic Capital ensures that individuals have broadly distributed equity holdings of AI companies. By “predistributing” AI gains through ownership rather than relying on ex-post redistribution, these mechanisms address inequality at its source and ensure that all citizens have a stake in AI advancement.

Equity-based approaches insure against the uncertainty inherent in technological progress, both regarding the timing of AI advances and regarding how transformative it will be. If AI development stalls, returns remain modest; if AI transforms the economy, returns are likely to rise. This automatic adjustment proves valuable given radical uncertainty surrounding AI development. Moreover, the heavy reliance of our current tax system on labor taxation and relatively low effective tax rates on capital imply that equity-based approaches would offer insurance precisely when the labor share of the economy and by extension tax revenue decline

significantly due to transformative AI.

Such mechanisms suit stage 1 well, capturing value during labor's decline if consumption taxation faces constraints. They may also slow down the transition to stage 2 when much of the economy's value creation is driven by AI but does not significantly contribute to human consumption. Equity-based approaches may even serve as transitional mechanisms between current tax systems and the radical reforms that may be required for AGI-dominated economies. Early equity accumulation could provide an institutional basis for benefit-sharing as AI entities evolve toward more and more autonomous AGI.

5 Conclusion

This paper offers concrete tax policy lessons for what may be the most transformative economic shift in human history. Transformative AI may fundamentally reshape the economy. Our analysis explores how tax systems can be adapted to harness this transformation for shared prosperity.

Our investigation considers this transition in two steps. In stage 1, if AI reduces labor's importance in production, consumption taxation emerges as the primary tool for revenue generation and redistribution. Differential commodity taxation, which is widely regarded as suboptimal in today's economy, may regain importance as labor distortions diminish. The declining importance of labor would remove the primary constraint on taxation, raising the returns to more sophisticated tax designs that better address distortions arising from evasion, household production, and other factors.

Stage 2 looks further ahead to the AGI era, where artificial systems might capture the majority of economic surplus. In our analysis, the planner's problem reduces to familiar economic principles, akin to an optimal harvesting problem where society must decide how much to "harvest" AGI's growing capital stock for human benefit. We find that the optimal tax rate on AGI in such scenarios depends crucially on human time preferences. Hence, human values and collective choices play a central role in shaping optimal institutions.

Given the magnitude of the potential challenges for public finance, we expect the path forward to include proactive institutional adaptation, beginning with reforms to consumption taxation systems while building capacity for the eventual taxation of AGI capital. Our analysis

finds that a shift from labor- toward first consumption-based and then AGI-capital-based taxation may help maintain fiscal sustainability, equity, and efficiency. Strengthening consumption tax infrastructure today could prepare the way for stage 1 of the transition. Similarly, existing corporate tax frameworks and emerging equity-based mechanisms (such as sovereign wealth funds and windfall clauses) may provide institutional foundations for eventually harvesting the fruits of AGI in stage 2. This gradual and deliberate transition may generate higher welfare for current and future generations than attempting radical reforms after disruption has already occurred.

Acknowledgments

This paper was originally prepared for the Brookings Center on Regulation and Markets (CRM) and presented at a Brookings CRM Authors' Conference. We are grateful to Bill Gale and Matthew Weinzierl for their insightful reviews, and to Era Dabla-Norris, Anh Nguyen, Phil Trammell, and participants at a Brookings CRM Authors' Conference and the NBER Economics of Transformative AI Workshop for helpful comments. Korinek created part of this work as a Visiting Fellow at the Brookings Institution. The views expressed in this work do not necessarily reflect the views of the Brookings Institution. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economics-transformative-ai/public-finance-age-ai-primer>.

References

- Acemoglu, Daron. 2024. "Harms of AI." In *The Oxford Handbook of AI Governance*, edited by Justin B. Bullock et al. Oxford University Press.
- Adam, Stuart. 2011. *Tax by Design: The Mirrlees Review*. Vol. 2. Oxford University Press.
- Altig, David, Alan J. Auerbach, Laurence J. Kotlikoff, Kent A. Smetters, and Jan Walliser. 2001. "Simulating Fundamental Tax Reform in the United States." *American Economic Review* 91 (3): 574–95.
- Amodei, Dario. 2024. "Machines of Loving Grace." <https://darioamodei.com/machines-of-loving-grace>.
- Atkinson, Anthony B., and Joseph E. Stiglitz. 1976. "The Design of Tax Structure: Direct versus

- Indirect Taxation." *Journal of Public Economics* 6 (1–2): 55–75.
- Auerbach, Alan J., and Laurence J. Kotlikoff. 1987. "Evaluating Fiscal Policy with a Dynamic Simulation Model." *The American Economic Review* 77 (2): 49–55.
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Brynjolfsson, Erik, Anton Korinek, and Ajay Agrawal. 2025. "The Economics of Transformative AI: A Research Agenda." forthcoming.
- Chamley, Christophe. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica* 54 (3): 607–22.
- Congressional Budget Office. 2014. "Taxing Capital Income: Effective Marginal Tax Rates under 2014 Law and Selected Policy Options." CBO Report 49817. Congressional Budget Office, Washington, DC. Accessed December 8, 2025.
- Corlett, W. J., and D. C. Hague. 1953. "Complementarity and the Excess Burden of Taxation." *The Review of Economic Studies* 21 (1): 21–30.
- Diamond, Peter A., and James A. Mirrlees. 1971. "Optimal Taxation and Public Production I: Production Efficiency." *The American Economic Review* 61 (1): 8–27.
- Diamond, Peter, and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives* 25 (4): 165–90.
- Gale, William G. 2020. "Raising Revenue with a Progressive Value-Added Tax." In *Tackling the Tax Code: Efficient and Equitable Ways to Raise Revenue*, edited by Emily Moss, Ryan Nunn, and Jay Shambaugh. The Hamilton Project.
- George, Henry. 1879. *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions and of Increase of Want with Increase of Wealth: The Remedy*. D. Appleton and Company, New York.
- Huynh, Ky-Cuong, Suchet Mittal, and Noah Frank. 2025. "Funding Government in the Age of AI." *Convergence Analysis*.
- <https://www.convergenceanalysis.org/fellowships/spar-economics/funding-government-in-the-age-of-ai>.
- Institute for Fiscal Studies. 1978. *The Structure and Reform of Direct Taxation*. George Allen & Unwin, London.

- Jones, Charles I. 2024. "The AI Dilemma: Growth versus Existential Risk." *American Economic Review: Insights* 6 (4): 575–90.
- Judd, Kenneth L. 1985. "Redistributive Taxation in a Simple Perfect Foresight Model." *Journal of Public Economics* 28 (1): 59–83.
- Kaplow, Louis. 2010. "Taxing Leisure Complements." *Economic Inquiry* 48 (4): 1065–71.
- Karnofsky, Holden. 2021. "Forecasting Transformative AI, Part 1: What Kind of AI?" *Cold Takes* (blog), August 10, 2021.
- Kleven, Henrik Jacobsen. 2004. "Optimum Taxation and the Allocation of Time." *Journal of Public Economics* 88 (3–4): 545–57.
- Korinek, Anton, and Megan Juelfs. 2024. "Preparing for the (Non-Existent?) Future of Work." In *Oxford Handbook of AI Governance*, edited by Justin Bullock et al., 746–76. Oxford University Press.
- Korinek, Anton, and Joseph E. Stiglitz. 2019. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. University of Chicago Press.
- Korinek, Anton, and Joseph E. Stiglitz. 2025. "Steering Technological Progress." *INET Working Paper*.
- Korinek, Anton, and Donghyun Suh. 2024. "Scenario for the Transition to AGI." NBER Working Paper No. 32255. Cambridge, MA: National Bureau of Economic Research.
- Korinek, Anton, and Jai Vipra. 2025. "Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence." *Economic Policy* 40 (121): 225–56.
- Mankiw, N. Gregory, Matthew Weinzierl, and Danny Yagan. 2009. "Optimal Taxation in Theory and Practice." *Journal of Economic Perspectives* 23 (4): 147–74.
- Mirrlees, James A., and Stuart Adam. 2010. *Dimensions of Tax Design: The Mirrlees Review*. Oxford University Press.
- OpenAI. 2018. "OpenAI Charter." <https://openai.com/charter/>. Accessed December 8, 2025.
- Ramsey, Frank P. 1928. "A Mathematical Theory of Saving." *The Economic Journal* 38 (152): 543–59.
- Stevenson, Betsey. 2025. "What Is There to Fear in a Post AGI World." In *The Economics of*

Transformative AI, edited by Ajay Agrawal, Erik Brynjolfsson, and Anton Korinek. University of Chicago Press. This volume.

Susskind, Daniel. 2020. *A World Without Work: Technology, Automation, and How We Should Respond*. Metropolitan Books, New York.

Tegmark, Max. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf, New York.